



22883

PATENT TRADEMARK OFFICE

136.1005.01

This application is submitted in the name of the following inventors:

| <u>Inventor</u> | <u>Citizenship</u> | <u>Residence City and State</u> |
|-------------------|--------------------|---------------------------------|
| Jack Regula | USA | San Jose, CA |
| Jhy-Ping Shaw | USA | San Jose, CA |
| Ronald A. Simmons | USA | Bountiful, UT |
| Curtis Winward | USA | Lehi, UT |
| Ralph Woodard | USA | Mountain View, CA |
| William Wu | USA | Cupertino, CA |

The assignee is PLX Technology, Inc., a corporation having an office at 870 Maude Ave., Sunnyvale, CA 94085.

Title of the Invention

On-Chip Switch Fabric

Background of the Invention

Field of the Invention

This invention relates to on-chip communication. In particular, the invention relates to on-chip communication between plural on-chip components across a communication bus having plural tracks.

Description of the Related Art

Traditional integrated circuit chips, also known as computer chips, are dedicated to a single function, with the chips attached to one another at a circuit board level. However, the number and types of circuits that can be place on a computer chip has continued to advance at a rapid pace. It is now possible to include circuits for many different functions on a single chip to create a complete “system on a chip.”

Designing systems on a chip can be daunting. In particular, providing for communication between different on-chip integrated components can be difficult. Furthermore, traditional design approaches tend not to be scalable to systems that involved increasing numbers of on-chip components.

Each function on multi-functional single chip is implemented by an independently operating module. To function, each module exchanges data with another module. These modules function as a data transfer pair. As the number of functions on a single chip increases, multiple data transfer pairs are needed to simultaneously transfer data. In a traditional time domain shared bus, only one data transfer pair can transfer data on the shared bus at any given time. Thus, in the event that multiple data transfer pairs need to simultaneously transfer data, only one pair can have access to the bus at a time and the other pairs must wait. In a switch fabric, each module has a communication path from itself to all other modules; and thus, if the target module is not currently engaged in a data transfer, it can accept data from an initiator without contention with other data transfers that may be simultaneously occurring.

Summary of the Invention

Accordingly, what is needed is a system for providing simultaneous communication among on-chip integrated components. This system should be flexible enough to accommodate different types of components. The system also should allow for easy integration of the components. Furthermore, the system should be easily scalable – in terms of both bandwidth and connectivity - to provide communication between increasing numbers of integrated components.

The invention addresses the foregoing needs by providing a system that includes an on-chip communication switch fabric for use by on-chip components. Preferably, the system uses a zero-wait-state packet-based communication protocol. The primary reason for packet based data transfers is because any target may have multiple initiators desiring to transfer data to it at any given time. By using a zero-wait-state packet based data transfer, the initiator is forced to transfer data every clock cycle which maximizes the data transfer bandwidth to the target. A second reason is that by limiting the packet size, the arbiter must frequently re-arbit and grant the bus; this ensures that the bus will operate in accordance with the priority scheme that the arbiter is designed to implement. Each target and initiator have exactly the same interface signals and timing, greatly simplifying learning costs for chip developers. The system also preferably uses multiplexors for signal selection, with the multiplexors being constructed from plural smaller multiplexors that can be distributed across a chip. This feature allows the system to be spread out across a chip, facilitating scalability. Furthermore, in a preferred embodiment, the system can use a different clock domain from the components, allowing for greater flexibility in chip design. Each component, as well as the system, may be in an independent clock domain.

1 Accordingly, one embodiment of the invention is a system for communication on
2 a chip. The system includes an on-chip communication bus including plural tracks, and a
3 plurality of stations that couple a plurality of on-chip components to the on-chip communication
4 bus. The plurality of on-chip components use the tracks to communicate. Preferably, the stations
5 use a packet based communication protocol. Each component has a dedicated track which it can
6 use to send information to any/all other components.

7
8 Example of on-chip components that can utilize the invention include, but are not
9 limited to, a PCI bridge, a USB interface, and an I2C interface. Other examples include a UART
10 interface, a DDR and/or SDRAM, an ethernet interface, a general I/O interface, and other
11 components.

12
13 In a preferred embodiment, each station includes an initiator that requests
14 permission to transmit outgoing data over a track to another station and that transmits the
15 outgoing data, an arbiter that evaluates requests from other stations and selects a track on which
16 to receive incoming data, and a target that receives the incoming data. The arbiter is constructed
17 to receive requests of varying priorities and to grant access based upon those priorities. The
18 initiator can be connected to a grant multiplexor for selecting a grant line, and the arbiter can be
19 connected to a track multiplexor for selecting a track. In order to facilitate scalability, these
20 multiplexors can be constructed from plural smaller multiplexors distributed across the chip.
21 The plurality of tracks and multiplexors preferably implement a crossbar switch.

22
23 Each station can also include a source queue for queuing outgoing data and a
24 destination queue for queuing incoming data. These queues preferably are first-in-first-out

1 registers. The source queue and the destination queue can serve to separate a clock domain for
2 the on-chip communication bus from clock domains for the plurality of on-chip components.
3 Thus, components that run at different clock speeds can be more easily accommodated than in
4 traditional systems.

5
6 In order to provide for even greater flexibility, more than one of the plurality of
7 on-chip components can be coupled to the on-chip communication bus through one of the
8 stations. This arrangement is particularly useful for connecting plural slower components to the
9 bus, with the benefit that memory and routing resources can be conserved.

10
11 Each station also preferably includes or is connected to a watchdog circuit that
12 determines if its station has gone offline. If a watchdog station determines that its station has
13 gone offline, that watchdog station informs a controller connected to the system. The controller
14 can then re-route or block communications to that station, thereby helping to prevent the offline
15 station from interfering with normal communications between components across the system.
16

17 The invention also includes methods for performing the foregoing operations, as
18 well as other embodiments of the invention.
19

20 This brief summary has been provided so that the nature of the invention may be
21 understood quickly. A more complete understanding of the invention may be obtained by
22 reference to the following description of the preferred embodiments thereof in connection with
23 the attached drawings.
24

Brief Description of the Drawings

Figure 1 shows an overview of an on-chip communication system according to the invention.

Figure 2 illustrates one possible embodiment of an on-chip communication system according to the invention.

Figure 3 illustrates one possible embodiment of a station for an on-chip communication system according to the invention.

Figure 4 illustrates one possible arrangement for interconnecting track lines for stations according to the invention.

Figure 5 illustrates one possible arrangement for interconnecting grant lines for stations according to the invention.

Figure 6 illustrates one possible arrangement for plural on-chip components to share a single station according to the invention.

Figure 7 is a flowchart for explaining communication between components across an on-chip communication system according to the invention.

Figure 8 illustrates a technique for interconnecting stations using smaller multiplexors to improve scalability according to the invention.

Description of the Preferred Embodiment

Related Disclosure

Inventions described in this disclosure can be used in conjunction with inventions described in the following application: Application Serial No. _____, filed _____, in the name of inventor Jack Regula, for Ordering in Multi-Path Switching Fabrics, Express Mail No. _____. This application is hereby incorporated by reference as if fully set forth herein.

Lexicography

Chip: An integrated circuit chip. Examples include, but are not limited to, a central processing unit, digital signal processing chip, memory manager, or complete "system-on-a-chip."

System-on-a-Chip: A chip that contains all circuits necessary for implementing a complete system, for example for a basic computer.

Component: A subset of circuits on a chip that perform a particular function or operation. Examples include, but are not limited to, a PCI (peripheral component interconnect)

1 bridge, a USB (universal serial bus) interface, an I2C (inter-integrated-circuit) interface, a UART
2 (universal asynchronous receiver transmitter) interface, a DDR (data direction register) and/or
3 SDRAM (synchronous dynamic access memory), an ethernet interface, a general I/O
4 (input/output) interface, and other circuits and interfaces. Components also can be referred to as
5 peripherals.

6
7 Station: A port to an on-chip communication bus according to the invention.

8
9 Clock Domain: A subset of circuits or components that uses a common clock
10 signal.

11
12 Packet-Based Protocol: A communication protocol in which data is sent in
13 packets, typically along with header information for the data.

14
15 Split-Response Transaction: A two-stage operation that is split over two
16 transactions, namely a request operation and a completion operation. In a split-response read
17 transaction, a first station sends a read request to a second station. The second station responds
18 to the read request command by initiating a read completion operation to write the requested data
19 to the first station.

20
21 Head-of-Line Blocking: Blocking that occurs when transmission of data at the
22 front of a source queue is delayed because it is intended for a station or component that is busy,
23 thereby blocking transmission of data deeper in the source queue that is intended for a station or
24 component that is not busy.

1 Queue: A register or memory that stores data while the data awaits transmission
2 or other processing.

3
4 FIFO (First In First Out) Register: A register that orders data such that data is sent
5 from the register in the order that the data was received by the register.

6
7 *Overview*

8
9 Figure 1 shows an overview of an on-chip communication system according to the
10 invention.

11
12 Chip 1 in Figure 1 includes plural on-chip components that communicate using an
13 on-chip communication system. The components in Figure 1 are PCI bridge 2, USB interface 3,
14 UART interface 4, I2C interface 5, DDR and SDRAM 6, EEPROM 7, Ethernet interface 8,
15 general I/O interface 9, and other components 10 and 11. Each of these components is connected
16 to on-chip communication bus 12 through stations 13 to 22, respectively. Thus, components 2 to
17 11 can communicate with each other through stations 13 to 22 and on-chip communication bus
18 12. The invention is not limited to the particular number and/or types of components shown in
19 Figure 1.

20
21 According to the invention, on-chip communication bus 12 includes plural tracks.
22 These plural tracks allow more than one component to communicate with another component
23 simultaneously.

Each track preferably includes lines for data bits and other control information. For example, one embodiment of a track includes lines for 64 bits of data, eight command/byte enable (C/BE) signals, two parity signals (one per double word of data), a start of packet signal, and an end of packet signal.

On-chip communication bus 12 preferably uses a packet based communication protocol. Use of such a protocol simplifies a chip designer's task in developing and/or modifying components to communicate through the on-chip communication bus and reduces the time that an initiator consumes for a given size data transfer. The underlying principle is that a station does not initiate a data transfer until it is ready to communicate quickly.

Station Design and Interconnection

Figure 2 illustrates one possible embodiment of an on-chip communication system according to the invention. Figure 2 is a high-level diagram that shows the basic functionality used by stations according to the invention.

Briefly, a system for communication on a chip includes an on-chip communication bus including plural tracks, and a plurality of stations that couple a plurality of on-chip components to the on-chip communication bus. Each station has a dedicated track which it can use to send information to other stations.

1 In Figure 2, stations 25 to 28 intercommunicate through switch fabric 29, which
2 includes on-chip communication bus 12. Of course, the invention is not limited to four stations,
3 and the stations need not be constructed and arranged as shown in Figure 2.
4

5 Each of stations 25 to 28 is constructed similarly. Station A 25 includes
6 transmitter 31, requester 32, receiver 33 and arbiter 34. Station B 26 includes transmitter 36,
7 requester 37, receiver 38 and arbiter 39. Station C 27 includes transmitter 41, requester 42,
8 receiver 43 and arbiter 44. Station D 28 includes transmitter 46, requester 47, receiver 48 and
9 arbiter 49. While the transmitters, requesters, receivers and arbiters are shown as separate blocks
10 in Figure 2, these functions can be combined in a single circuit or block.
11

12 Transmitter 31 of station A 25 is responsible for transmitting data to switch fabric
13 29. In Figure 2, clocking of data from transmitter 31 is enabled by requester 32 through a clock
14 enable (CLKEN) signal.
15

16 Before requester 32 of station A 25 enables transmission of data, requester 32
17 sends a request (REQ) signal to each of the other stations connected to switch fabric 29. In a
18 preferred embodiment of the invention, the request signals are multi-bit signals that incorporate
19 different levels of priority for requests. For example, in a preferred embodiment, each request
20 line is three bits wide to allow for seven different request priority levels (plus a no-request level
21 of 000). When requester 32 receives a grant (GNT) signal from one of the other station in
22 response to the request signal, requester 32 enables transmission of data from transmitter 31.
23

1 Station A 33 also can receive data, in particular through receiver 33. Arbiter 34 of
2 station A 25 arbitrates and controls what data is sent to station A 25 from the other stations.
3 Arbiter 34 performs this arbitration based on the priorities of request signals sent from the other
4 stations. Arbiter 34 controls what data is sent to station A 25 by sending various grant signals in
5 response to those request signals. This arrangement, in which a station can select what data is
6 sent to that station, allows implementation of a split-response transaction model for
7 communication over switch fabric 29.

8
9 Stations B 26 to D 28 operate similarly to station A 25.

10 The components connected each of the stations are not shown in Figure 2. These
11 components provide the data sent by the transmitters and receive the data received by the
12 receivers. One or more such components can be connected to each station.
13

14 Other elements also can be included in the on-chip communication system
15 according to the invention. For example, the system can include system registers for storing
16 system parameters and a system controller for controlling system operation. These system
17 registers and system controller preferably are connected to the on-chip communication system
18 through their own station. The system also can include other special stations, watchdog circuits,
19 and other elements.
20

21
22 Figure 3 illustrates a preferred embodiment of a station for an on-chip
23 communication system according to the invention.
24

1 In Figure 3, station 50 connects component 51 to switch fabric 52, which includes
2 on-chip communication bus 53 with plural tracks. Thus, in order for component 51 to
3 communicate with other components across on-chip communication bus 53, component 51
4 transmits data to and receives data from station 50. Station 50 in turn communicates with other
5 stations through switch fabric 52, and those stations communicate with their respective
6 components.

7
8 Station 50 preferably includes initiator 54, target arbiter 55, and target 56.
9 Initiator 54 requests permission to transmit outgoing data over a track to another station and
10 transmits the outgoing data. Target arbiter 55 evaluates requests from other stations and selects a
11 track on which to receive incoming data. Target 56 receives the incoming data.

12
13 Compared to the stations shown in Figure 2, initiator 54 performs the functions of
14 both a requester and a transmitter shown in Figure 2. Target arbiter 55 performs the functions of
15 an arbiter shown in Figure 2. Target 56 performs the functions of a receiver shown in Figure 2.

16
17 Returning to Figure 3, initiator 54 is connected to multiplexor 57, which in turn is
18 connected to on-chip communication bus 53. Likewise, target 56 is connected to multiplexor 58,
19 which also is connected to on-chip communication bus 53.

20
21 The multiplexors for all stations connected to the on-chip communication bus
22 along with the tracks of the bus form switch fabric 52, which preferably implements a crossbar
23 switch. The switch fabric also can include other elements, as discussed in more detail with
24 respect to Figure 8. This switch fabric serves to switch data between stations over the tracks of

1 on-chip communication bus 53. The switch fabric also can switch grant signals and other control
2 data.

3
4 In order for the invention to utilize the plural tracks of on-chip communication
5 bus 53, switch fabric 52 preferably is a multi-path switch fabric. In a preferred embodiment, this
6 multi-path switch fabric is substantially equivalent to a cross-bar switch, except that the
7 invention preferably utilizes arbitration based on request signals to determine switching as
8 opposed to conventional scheduling.

9
10 Initiator 54 in Figure 2 also is connected to source queue 60, and target arbiter 55
11 and target 56 are connected to destination queue 61. These queues preferably are first-in-first-out
12 (FIFO) registers.

13
14 Queues 61 and 62 allow component 51 to operate in a different clock domain (i.e.,
15 using a different clock speed and/or clock) from the on-chip communication bus, and thus in a
16 different clock domain from other components. Figure 3 shows on-chip communication bus
17 clock domain 63 on one side of queues 60 and 61, and component clock domain 64 on the other
18 side of queues 60 and 61.

19
20 Different clock domains can be accommodated because data can be clocked into
21 the queues at a different rate than the data is clocked out. This provides chip designers with
22 greater flexibility in designing chips and integrating different components into those chips as
23 compared to systems in which only one or a few clock domains can be accommodated.

1 Of course, the invention does not require that components run in different clock
2 domains. Components can run in the same clock domain as the on-chip communication bus
3 and/or each other, if so desired.

4
5 Source queue 60 is connected to packetizer 66, and destination queue 61 is
6 connected to de-packetizer 67. The packetizer and de-packetizer allow component 51 to
7 communicate with station 50 using a simplified packet-based protocol. Use of such a protocol
8 simplifies the task of connecting a component to a station according to the invention, thereby
9 reducing learning costs for chip designers using the invention.

10
11 A preferred embodiment of the packet protocol uses a 64 bit header and variable-
12 sized payloads. Up to 32 payloads preferably can be sent with each header. The preferred
13 embodiment of the header includes the following fields: station ID, report bit, long address bit,
14 priority field, tag field, payload count, and address.

15
16 The station ID is 5 bits and identifies the source of the packet. It is assigned by
17 the chip designer.

18
19 The report bit indicates whether or not a destination station should report to a
20 source station with a completed without error message after completion of a data transfer or other
21 command without an error.

22
23 The long address bit indicates that the first 24 bits of the first payload after the
24 header contains additional address information.

1 The priority field holds a 3 bit priority level for the packet. This priority
2 preferably matches the priority of the request signal sent for the packet.

3
4 The tag field is a 5 bit field used to uniquely identify split-response transaction
5 requests. These types of requests are used in read operation, as discussed in more detail below
6 with reference to Figure 7.

7
8 The payload count contains 9 bits that indicate how many packets of payload are
9 associated with and will follow the header.

10
11 The address field stores a 40 bit address for the data. This address preferably is
12 with respect to an address space assigned to the station, and thereby to the component(s)
13 connected to the station.

14
15 Other arrangements for a station, component, switch fabric and packet layout are
16 possible and also fall within the scope of the invention.

17
18 Figure 4 illustrates one possible arrangement for interconnecting track lines for
19 stations according to the invention.

20
21 In Figure 4, initiators 69 to 72 are connected to targets 74 to 77 through
22 multiplexors 79 to 82. The initiators, targets and multiplexors are connected such that data sent
23 over a track from an initiator at any station can be received by a target at any other station. The
24 multiplexors in Figure 4 correspond to multiplexor 58 in Figure 3. Thus, when a station's target

1 arbiter sends a grant signal to another station, that target arbiter uses the station's track
2 multiplexor to select the corresponding track for receiving data from the other station.

3
4 Figure 5 illustrates one possible arrangement for interconnecting grant lines for
5 stations according to the invention.

6
7 In Figure 5, target arbiters 84 to 87 are connected to initiators 89 to 92 through
8 multiplexors 84 to 97. The target arbiters, initiators and multiplexors are connected such that
9 grant signals sent from an arbiter at any station can be received by an initiator at any other
10 station. The multiplexors in Figure 5 correspond to multiplexor 57 in Figure 3. Thus, when a
11 station's initiator sends a request signal to another station, that initiator uses the station's grant
12 multiplexor to select the corresponding grant line from the target arbiter for the other station.
13 The requesting station can then monitor that grant line for a grant signal from the other station.

14
15 The request lines preferably are not connected to the stations through
16 multiplexors. Instead, the request line(s) from each station's initiator preferably are directly
17 connected to each other station's target arbiter. Each station's target arbiter preferably is directly
18 connected to all request lines from all other stations. For example, if there are four stations, each
19 station's target arbiter preferably is connected to the three sets of request lines from each of the
20 other stations. This arrangement allows stations to receive and to react extremely quickly to
21 request signals from other stations.

Station Sharing

Figure 6 illustrates one possible arrangement for plural on-chip components to share a single station according to the invention. This arrangement is particularly useful when several components are relatively slow compared to other components and/or to the on-chip communication system.

In Figure 6, three components 100 to 102 share a station. These components are illustrated as USB interface 100, UART interface 101, and I2C interface 102. Of course, the invention is not limited to these particular components or to three components sharing a station. More or fewer components can share a station according to the invention.

As shown in Figure 6, an additional arbiter 104, decoder 105 and multiplexor 106 are used to connect the plural components to a station. Comparing Figures 3 and 6, all of the elements in Figure 6 take the place of component 51 in Figure 3, with additional signals provided for address and flag information.

Arbiter 104 in Figure 6 further arbitrates grants and requests among the sharing components. Decoder 105 decodes address and flag information so as to route incoming data to the appropriate component. Multiplexor 106 likewise selects outgoing data from the appropriate one of the components.

1 Other arrangements for sharing a station are possible and also fall within the scope
2 of the invention. In any case, sharing of a station by plural components conserves memory and
3 routing resources.

5 *Split-Response Transaction Model*

7 The invention utilizes a split-response transaction model of communication. A
8 write operation from one station to another is simple in this model. A first station requests
9 permission to write to a second station. If the second station is available and has room in its
10 incoming packet buffer, the second station grants the request. Then, the first station sends a write
11 command to the second station, followed by the data.

13 A read operation is slightly more complicated because a station preferably needs
14 to make data available before it can be returned to the requesting station. In order to perform a
15 read operation, a first station again requests permission to send a read request to a second station.
16 However, instead of sending data, the first station sends a read request command. This command
17 preferably includes address information for the data to be read.

19 The second station responds to the read request command by initiating a read
20 completion operation to write the requested data to the first station. This read completion
21 operation is substantially identical to a write operation from the second station to the first station,
22 except that the second station indicates that the operation is a read completion. The second
23 station preferably makes this indication through the bus command portion of the track used to
24 send the data for the read operation.

1
2 The two-stage read operation is called a “split-response transaction” operation
3 because the operation is split over two transactions: a read request and a read completion. The
4 tag field in the header for any packets sent in response to a read request is used to align those
5 packets with the read request. In other words, the tag field is used to align a read request and the
6 resulting data across the split-response transaction.

7
8 Using the foregoing approach, all operations between stations involve transmitting
9 information from one station to another station for consumption. For a write, the information
10 includes a write command and the actual data to be written. For a read request, the information
11 includes a read request command and address information. For a read completion, the
12 information includes the data that was requested by the corresponding read request command,
13 along with an indicator that the data is for a read completion command.

14 15 *Transmitting Information* 16

17 Figure 7 is a flowchart for explaining communication between components across
18 an on-chip communication system according to the invention. The steps in Figure 7 are
19 discussed with reference to the elements of the station depicted in Figure 3 in order to improve
20 understanding of the invention. However, the method illustrated by Figure 7 is not limited to use
21 with the station shown in Figure 3.

22
23 In step S701, a component communicates with its station to request a data transfer
24 over the on-chip communication bus with another component connected to another station. This

1 data transfer could be a write operation or a read operation. The first and second stations
2 communicate with each other to accomplish the data transfer in steps S702 to S709.

3
4 In step S702, the first station's initiator sends a request signal to the second
5 station in step S702. This request is received by the second station's target arbiter.

6
7 As discussed above, the request signal preferably indicates a priority for the data
8 transfer. Write operations preferably are assigned higher priorities than read operations. Thus,
9 when requests are evaluated by the second station's target arbiter, writes can be executed before
10 any pending reads. This priority scheme facilitates use of the split-response transaction model
11 for communication between components. Without this priority scheme, a station could choose to
12 read (i.e., consume) data before an earlier-issued write was completed, possibly causing the
13 station to inadvertently read stale or inaccurate data.

14
15 In step S703, the second station's target arbiter evaluates all outstanding requests
16 from other stations, including the request from the first station. The target arbiter preferably
17 selects the request with the highest priority.

18
19 In order to grant the first station's request, the second station's target arbiter sends
20 a grant signal to the first station in step S704. In step S705, the second station selects a track for
21 the data. In actual operation, steps S704 and S705 preferably occur simultaneously by sending a
22 grant signal from the second station's target arbiter to both the first station and to a track
23 multiplexor in the second station.

1 In response to the grant signal, the first station's initiator sends a command and/or
2 data to the second station in step S706. The command preferably is sent using the command/byte
3 enable signal lines of the selected track. Commands include, but are not limited to, write
4 commands, read request commands, and read completion commands. The data preferably is sent
5 using the 64 data lines in the selected track.
6

7 In step S707, the target at the second station receives the command and/or data.
8 Then, if the command is a read request, flow proceeds from step S707 through step S708 to step
9 S709. In step S709, the first and second stations reverse roles, and the station that received the
10 read request initiates a read completion command to send the data.
11

12 *Other Operations*

13

14 The on-chip communication system according to the invention also preferably can
15 execute register read and write operations for reading and writing to system registers. Because
16 these system registers preferably also are connected to the on-chip communication system
17 through a station, the process of reading and writing to the system registers is similar to that
18 discussed above. Additionally, the system preferably can execute special I/O commands, system
19 control commands (e.g., initialize, abort, etc.), and the like. System commands preferably are
20 directed toward a system controller connected to the system through a station.
21

22 The invention also can accommodate special direct memory access operations
23 among stations. These operations involve a special direct memory access station that is beyond

1 the scope of this disclosure. However, such stations can be connected to the on-chip
2 communication system disclosed herein without departing from the foregoing teachings.

3 *Head-of-Line Blocking*

4
5 Head-of-line blocking occurs when transmission of data at the front of a source
6 queue is delayed because it is intended for a station or component that is busy, thereby blocking
7 transmission of data deeper in the source queue that is intended for a station or component that is
8 not busy. This type of blocking can greatly impact communication in a system.

9
10 The invention addresses head-of-line blocking in at least three ways: through use
11 of “tracks” that typically have twice as much bandwidth as is required by the source or
12 destination of the data, through use of a packet-based communication protocol, and through use
13 of a watchdog circuit.

14
15 The on-chip communication system according to the invention can be very fast.
16 Thus, any blocking that occurs is not likely to last long. This strength is enhanced by the
17 system’s ability to use a different clock domain for the communication bus than the components
18 connected to the system. As a result, the on-chip communication system can operate at a higher
19 clock speed than the components, further reducing the impact of any blocking. It is well known
20 that head of line blocking limits throughput to roughly 59% of the peak speed of the interconnect.
21 By having a 2:1 overspeed in the interconnect, we allow sources and destinations to achieve their
22 full data rate despite head of line blocking.
23

1 The packet-based protocol used by the invention preferably limits how many
2 payloads and the length of each payload that can be sent in response to a grant of a request to
3 send data. As a result, no one data transfer operation is likely to tie up a station for too long,
4 thereby reducing the length of any blocking that does occur.

5
6 The on-chip communication system according to the invention also can include
7 one or more watchdog circuits. Preferably, one watchdog circuit is provided for each station.
8 These circuits can monitor the stations of the system to see if any station stalls or goes offline for
9 more than a predetermined amount of time (e.g., $\frac{1}{2}$ second). Preferably, the value for this
10 amount of time is stored in a system register for the on-chip communication system.

11
12 If a station stalls or goes offline for more than the predetermined amount of time,
13 that station's watchdog timer can inform a controller for the communication system. The
14 controller can then instruct all stations to purge any pending or queued operations involving the
15 offline or stalled station or to reroute those operations. Thus, if blocking occurs because of an
16 offline or stalled station, the blocking is terminated after the predetermined amount of time.

17 18 *Scalability*

19
20 The on-chip communication system according to the invention is scalable to large
21 systems. This scalability is possible because relatively few components are required to interface
22 each component to the system. Scalability also is facilitated by the ability of a station to interface
23 plural components to the system.

1 However, a problem does exist in that as the number of stations increases, the size
2 of the grant and track multiplexors also increases. This increase is not linear. Instead, the size of
3 the multiplexors increases by increasing amounts for each additional station. The increase is of
4 order N^2 , where N is the number of stations. At some point, if conventional multiplexor circuitry
5 is used, the footprint of the multiplexors on the chip can become too large and unwieldy to place
6 on the chip.

7
8 The invention addresses the foregoing issue by constructing the multiplexors from
9 smaller multiplexors and other circuits distributed across the chip. The stations are
10 interconnected using these smaller multiplexors, thereby alleviating the problem of having to
11 place large multiplexor circuits at one place for each station on the chip.

12
13 Figure 8 illustrates a technique for interconnecting stations using smaller
14 multiplexors to improve scalability according to the invention. The invention also includes the
15 use of pipeline storage elements – D flip-flops - between some of the multiplexor stages in order
16 to maintain transmission speed when a track must traverse a large number of multiplexor stages.
17 The invention also includes adjusting the time of issuance of a grant to a new transmitting station
18 relative to the end of transmission of a current transmitting station according to the number and
19 location of pipeline storage elements traversed by the track in the switch fabric in order to
20 eliminate idle cycles between the end of a transmission and the start of a next, waiting
21 transmission.

22
23 In Figure 8, transmitters/receivers 108 to 111 are interconnected through a switch
24 fabric including D flip-flops 113 to 117 and small multiplexors 119 to 126. In this case, the

1 term “small” is in comparison to larger multiplexors that would be needed using conventional
2 circuitry.

3
4 The dashed lines in Figure 8 illustrate connections that could be made to
5 accommodate more stations. Advantageously, no additional space need be used near the existing
6 stations. As a result, scalability is improved.

7
8 *Alternative Embodiments*

9 Although preferred embodiments of the invention are disclosed herein, many
10 variations are possible which remain within the content, scope and spirit of the invention, and
11 these variations would become clear to those skilled in the art after perusal of this application.